Partial Preferences for Mediated Bargaining

Piero A. Bonatti Marco Faella Luigi Sauro

Dept. of Electrical Engineering and Information Technologies
Università di Napoli "Federico II"
Italy

In this work we generalize standard Decision Theory by assuming that two outcomes can also be incomparable. Two motivating scenarios show how incomparability may be helpful to represent those situations where, due to lack of information, the decision maker would like to maintain different options *alive* and defer the final decision. In particular, a new axiomatization is given which turns out to be a weakening of the classical set of axioms used in Decision Theory. Preliminary results show how preferences involving complex distributions are related to judgments on single alternatives.

1 Introduction

In his pioneering work on Decision Theory [4], when delineating the fundamental properties of a preference relation $\prec$, Savage makes the following point: given two potential outcomes f and g, it cannot be the case that $f \prec g$ and $g \prec f$ at the same time. Clearly, this is logically equivalent to saying that either $f \not\prec g$ or $g \not\prec f$, which leads to three possible cases: *(i)* $f \not\prec g$ and $g \prec f$, *(ii)* $f \prec g$ and $g \not\prec f$, or *(iii)* $f \not\prec g$ and $g \not\prec f$. Then, he postulates that these three cases are the only possible judgments concerning f and g. In particular, the last case ($f \not\prec g$ and $g \not\prec f$) allegedly implies that f and g are equivalent in the sense that in any situation wherein these are the only two possible options, the decision maker does not mind delegating to coin flipping. Consequently, in classical Decision Theory (*CDT*) a very fundamental property of a preference relation is its totality.

From the theory's very start, the hidden assumptions underlying this model of an *economic man* raised some criticisms, one of the most influential of which being due to Simon:

> "This man is assumed to have knowledge of the relevant aspects of his environment which, if not absolutely complete, is at least impressively clear and voluminous. He is assumed also to have a well-organized and stable system of preferences, and a skill in computation that enables him to calculate, for the alternative courses of action that are available to him, which of these will permit him to reach the highest attainable point on his preference scale" [5].

In recent years, the massive development of e-commerce services makes Simon's criticisms even more cogent and the classical viewpoint on the economic man more and more idealistic. Often preferences result from complex trade-offs between different attributes (functionalities, cost, Quality of Service (QoS), information disclosure risks, etc.) sometimes the user has only a vague idea of. Moreover, in some cases the user actually consists of a group of persons where internal debate does not easily end up with a total preference. Finally, from a computer-science perspective, our aim could be to develop a software agent that acts in an electronic market on behalf of a real user. As we discuss in Section 2, even if the user conforms with the classical economic man, her preference relation could be so complex that it could not be entirely and efficiently injected into the software agent.

Differently from Simon, who moved towards a problem-solving perspective, in this work we challenge CDT on its playground. In particular, we provide an alternative axiomatization where two outcomes, due to the lack of information or an irreducible heterogeneity of the attributes involved, can also

F. Mogavero, A. Murano, and M.Y. Vardi (Eds.):
2nd Workshop on Strategic Reasoning 2014 (SR14)
EPTCS 146, 2014, pp. 105–111, doi:10.4204/EPTCS.146.14

© P.A. Bonatti, M. Faella & L. Sauro
This work is licensed under the
Creative Commons Attribution License.

be incomparable. In Section 2, we show two motivating scenarios where considering preferences as incomplete seems to be appropriate. In Section 3, we introduce the new axiomatization by emphasizing which axioms of CDT remain the same and which axioms should be replaced. The resulting theory, that we call Partial Decision Theory, consists in a weakening of CDT, that is, all the properties it satisfies are satisfied by CDT as well (but not vice versa). In Section 3, we show some general features of the new axiomatization; in particular, we argue that the proposed framework is not too weak, as it retains several desirable properties of CDT. Conclusions and future works end the paper.

2 Motivating Scenarios

In this section we introduce two scenarios where partial preferences seem to provide a more natural way to describe a decision maker, or a software agent behaving on its behalf, than total preferences.

In the first scenario, Bob wants to learn to play the piano and posts a request on a consumer-to-consumer social network. Soon after, he receives offers from two musicians, Carl and Mary. Carl provides two options: a 5 people class for 15 dollars per person or a one-to-one class for 35 dollars. Mary offers two similar options: a 3 people class for 20 dollars per person or a one-to-one class for 40 dollars. Furthermore, they both offer a trial lesson. Regarding Carl's options, Bob thinks that 5 people are too many for a class, thus he prefers the one-to-one option. On the contrary, he judges the price difference between Mary's options somewhat excessive, hence he prefers the 3 people class. If someone asks Bob *"Do you prefer Mary's 3 people class or Carl's one-to-one class?"*, Bob will probably answer *"I do not know, I first have to attend the trial lessons"*. Notice that this is different from saying that the two options are equivalent, because in that case Bob would simply flip a coin and choose one of them. On the contrary, it is more natural to think that these options are initially *incomparable* and Bob will use the trial lessons to disambiguate them.

More generally, in absence of complete information it might be difficult for an individual to figure out a coherent total order over the bids and choose in a single step one of them. On the contrary, making a decision can be viewed as a multiple step process where offers are initially filtered according to a partial preference relation. Then, depending on the resulting offers, an individual can acquire further information and possibly rank them.

In the second scenario, Alice's father has finally agreed to buy her a smartphone, and now she is browsing Ebay for possible offers. Unfortunately, the list is huge and patience is not Alice's forte. So, she would like to be assisted by a software agent to filter out undesired options. The software agent accepts constraints such as maximum cost and size, color restrictions, etc., and also a preference relation as a total order over bids. Then, according to the specified preference relation, the agent returns the best offer. Clearly, Alice's desires are influenced by several attributes such as operating system, color, weight, brand, and so on. For instance, she has a preference over operating systems in the following decreasing order: OS1, OS2 and OS3; over colors: blue, red, black, white; and over brands: Brand1, Brand2, Brand3. Furthermore, out of benevolence for her father, given a specific model the cheaper the better. However, such preferences over single attributes do not constitute a total order. Moreover, Alice cannot establish a priority over attributes, for instance she prefers a Brand3 phone with operating system OS1 to a Brand2 one with operating system OS2, but she prefers a Brand1 OS2 phone to a Brand2 OS1 one. Alice soon realizes that providing a total order to the software agent is frustrating and requires about the same effort as comparing all the offers by herself. This scenario reveals the following issue: in designing a software agent that behaves on behalf of real users, we have to take into account how users can *instruct* the agent about their own preferences. In electronic markets where the number of offers can be huge, it could be unfeasible to transfer an exact representation of users' desires into a software

agent. In this case, the agent should make do with an approximate representation of users' desires as a partial order and return a restricted list of choices from which the user can select the preferred one. As a further advantage, the user retains the ability of applying unforeseen, situation-specific knowledge and preferences that had not been formalized in advance.

3 Partial Preferences

Let $\Delta(\mathscr{A})$ be the class of all probability distributions over a countable set of alternatives $\mathscr{A}$. Given two probability distributions $f,g \in \Delta(\mathscr{A})$ and $\alpha \in [0,1]$, we denote by $\langle \alpha, f, g\rangle$ the convex combination of f and g such that $\langle \alpha, f, g\rangle(a) = \alpha f(a) + (1-\alpha)g(a)$, for all $a \in \mathscr{A}$. Moreover, for an alternative $a \in \mathscr{A}$, $[a]$ denotes the degenerate distribution that assigns probability 1 to a.

A preference relation $\preceq$ is a binary relation on $\Delta(\mathscr{A})$, subject to the following classical Decision Theory axioms: [1]

1. $f \preceq g$ or $g \preceq f$ *(totality)*;
2. if $f \preceq g$ and $g \preceq h$, then $f \preceq h$ *(transitivity)*;
3. if $f \prec g$ and $0 \leq \alpha < \beta \leq 1$, then $\langle \beta, f, g\rangle \prec \langle \alpha, f, g\rangle$;
4. if $f_1 \preceq g_1$, $f_2 \preceq g_2$, and $0 \leq \alpha \leq 1$, then $\langle \alpha, f_1, f_2\rangle \preceq \langle \alpha, g_1, g_2\rangle$;
5. if $f_1 \prec g_1$, $f_2 \preceq g_2$, and $0 \leq \alpha \leq 1$, then $\langle \alpha, f_1, f_2\rangle \prec \langle \alpha, g_1, g_2\rangle$;

where, as usual, $f \prec g$ means that $f \preceq g$ and $g \not\preceq f$.

Notice that the first two axioms force $\preceq$ to be a total (hence reflexive) transitive relation, i.e., a total preorder (also called non-strict weak order). As shown by the previous scenarios, we advocate that in several contexts some outcomes may be incomparable, meaning that $\preceq$ should be modeled as a (possibly partial) preorder. For this reason, we weaken the totality axiom in favor of one which requires reflexivity only:

1'. $f \preceq f$;

Having allowed for incomparable distributions, at first look it may seem that the deal is done. However, the obtained theory is so weak that it contemplates unrealistic preferences. The problem is that the previous axioms do not say anything about incomparable distributions which, once combined, can be then freely judged. On the contrary, it is natural to think that, to some extent, the incomparability between distributions persists also when they are combined.

Assume for example that f and g are incomparable and let $0 \leq \alpha < \beta \leq 1$. According to the axioms above, it is possible that $\langle \alpha, f, g\rangle \prec \langle \beta, f, g\rangle$. This looks inappropriate: given that I cannot compare f and g, why should I strictly prefer one combination of f and g over another? This leads to a further axiom:

6. if $0 \leq \alpha \leq 1$, and $\langle \alpha, f_1, f_2\rangle \prec \langle \alpha, g_1, g_2\rangle$, then there exist $j,k \in \{1,2\}$ such that $f_j \prec g_k$.

Intuitively, a distribution f of the type $\langle \alpha, f_1, f_2\rangle$ can be seen as a random choice (a.k.a. a *lottery*) which picks f_1 with probability α and f_2 with probability $1-\alpha$. Comparing f with another distribution g of the type $\langle \alpha, g_1, g_2\rangle$ encompasses comparing four possible draws: (f_1, g_1), (f_1, g_2), (f_2, g_1), and (f_2, g_2).[2] If there is no draw in which the second component is strictly better than the second, Axiom 6 requires that g is not strictly preferred to f. Notice that one could easily come up with more stringent

[1] Here, we borrow the formulation presented in [3].

[2] With probabilities α^2, $\alpha(1-\alpha)$, $\alpha(1-\alpha)$ and $(1-\alpha)^2$, respectively.

conditions on the persistence of incomparability, for instance by requiring that a majority of draws favors the second component over the first one. We instead propose a rather weak requirement, in the form of Axiom 6, which supports a wide range of preference relations, while still ensuring a number of interesting properties, which are the subject of Section 4.

We call *Partial Decision Theory* (PDT) the new set of axioms 1'-6. Notice that Axiom 6 can be easily derived in classical Decision Theory, consequently PDT is a weakening of classical Decision Theory, in the sense that all preference relations satisfying classical Decision Theory also satisfy PDT. The converse does not hold, as witnessed by the "empty" preference relation, i.e., the relation that considers incomparable all distinct distributions.

As seen above, Axiom 6 has been motivated by analyzing which preferences between $f = \langle \alpha, f_1, f_2 \rangle$ and $g = \langle \alpha, g_1, g_2 \rangle$ are admissible on the basis of the preferences on the possible draws (f_1, g_1), (f_1, g_2), (f_2, g_1), and (f_2, g_2). In Table 1 we perform such an analysis extensively. In particular, for each entry, the left-hand side $\bowtie_1 \bowtie_2 \bowtie_3 \bowtie_4$, with $\bowtie \in \{\sim, \prec, \succ, \lesssim\!\!\!\gtrsim\}$, is a consistent combination $f_1 \bowtie_1 g_1$, $f_1 \bowtie_2 g_2$, $f_2 \bowtie_3 g_1$, and $f_2 \bowtie_4 g_2$, whereas the right-hand side shows which preference relations between f and g are consistent with PDT. For example, the first entry is the case $f_1 \sim g_1$, $f_1 \sim g_2$, $f_2 \sim g_1$, and $f_2 \sim g_2$, then according with Axiom 4, $f \sim g$ is the only possibility. Conversely, in some other cases (e.g. $\prec \lesssim\!\!\!\gtrsim \succ \lesssim\!\!\!\gtrsim$) no axiom can be applied, consequently f can be in any relationship with g.

Table 1 provides a close look on how PDT behaves and hence it can be a good starting point to debate whether and how it can be extended or modified. For example, notice that in some cases f and g are comparable even if some of the underlying draws are not (e.g. due to Axiom 5, $\prec \lesssim\!\!\!\gtrsim \lesssim\!\!\!\gtrsim \sim$ results in $f \prec g$). Somewhat conversely, f and g can be incomparable even if all the underlying draws are comparable (e.g. $\prec \succ \prec \succ$ admits $f \lesssim\!\!\!\gtrsim g$). Finally, f and g are never forced to be incomparable, even if $f_1 \lesssim\!\!\!\gtrsim g_1$, $f_1 \lesssim\!\!\!\gtrsim g_2$, $f_2 \lesssim\!\!\!\gtrsim g_2$, and $f_2 \lesssim\!\!\!\gtrsim g_2$.

4 Properties of Partial Preferences

In the following, given two distributions f and g, we write $f \sim g$ for $f \preceq g$ and $g \preceq f$ (i.e., equivalence), and we write $f \lesssim\!\!\!\gtrsim g$ for $f \npreceq g$ and $g \npreceq f$ (i.e., incomparability). First, we show that two relevant properties of classical Decision Theory continue to hold in PDT. Given two distributions $f, g \in \Delta(\mathscr{A})$, we write $f \to g$ in case there exist $\varepsilon > 0$ and two alternatives a_1 and a_2 such that *(i)* $[a_1] \prec [a_2]$, *(ii)* $g(a_1) = f(a_1) - \varepsilon$, $g(a_2) = f(a_2) + \varepsilon$, and *(iii)* for all $a \neq a_1, a_2$, $g(a) = f(a)$. When $f \to g$, g can be obtained from f by shifting a positive amount of probability from an alternative a_1 to a strictly preferred alternative a_2. Then, denote by $f \Rightarrow g$ the transitive closure of $\to$.

Theorem 1 *Let $f, g \in \Delta(\mathscr{A})$. If $f \Rightarrow g$ then $f \prec g$.*

Proof. It suffices to show that $f \to g$ implies $f \prec g$. Assume that for some a_1 and a_2, where $[a_1] \prec [a_2]$, there exists $\varepsilon > 0$ such that $f(a_1) = g(a_1) + \varepsilon$ and $f(a_2) = g(a_2) - \varepsilon$. Let $\gamma = g(a_1) + g(a_2) = f(a_1) + f(a_2)$, $\alpha = \frac{f(a_2)}{\gamma}$ and $\beta = \frac{g(a_2)}{\gamma}$. Then, g can be written as $\langle \gamma, g', h \rangle$ and f as $\langle \gamma, f', h \rangle$, where h is a probability distribution such that $h(a_1) = h(a_2) = 0$, $g' = \langle \beta, [a_2], [a_1] \rangle$ and $f' = \langle \alpha, [a_2], [a_1] \rangle$. Since $[a_1] \prec [a_2]$ and $\alpha < \beta$, Axiom 3 implies that $f' \prec g'$. Then, due to Axiom 5, $f \prec g$. ∎

Another property that can be proved in PDT is that if each alternative coming out from a distribution f is dominated by all the alternatives from g, then $f \preceq g$. Preliminarily, given a distribution $f \in \Delta(\mathscr{A})$, the *support* of f, $supp(f) = \{a \in \mathscr{A} \mid f(a) > 0\}$, is the set of alternatives to which f assigns positive probability.

Theorem 2 *Let $f, g \in \Delta(\mathscr{A})$ be such that, for all $a \in supp(f)$ and $a' \in supp(g)$, $[a] \preceq [a']$. Then, $f \preceq g$.*

Table 1: Extensive analysis of PDT

Proof. We first show that for all $a \in supp(f)$, $[a] \preceq g$. The proof is by induction on the cardinality n of $supp(g)$ where the case $n = 1$ is trivial. Assume $n > 1$, then g can be written as $\langle \alpha, [a'], g' \rangle$, where a' is a generic alternative from $supp(g)$ and $g'(a') = 0$. Clearly, the distribution $[a]$ can also be written as $\langle \alpha, [a], [a] \rangle$. By assumption $a \preceq a'$ and, since the cardinality of g' is $n-1$, by induction $[a] \preceq g$. Then, by applying Axiom 4 we have $[a] \preceq g$.

Now, the proof is by induction on the cardinality m of $supp(f)$ where the base case $m = 1$ has been proved above. Assume $m > 1$ and let $a \in supp(f)$, then f can be written as $\langle \alpha, [a], f_{-a} \rangle$, where $f_{-a}(a) = 0$ and $f_{-a}(a') = \frac{f(a')}{1-f(a)}$ for all $a' \in supp(f) \setminus \{a\}$. We already proved that $[a] \preceq g$ and by induction hypothesis it holds also that $f_{-a} \preceq g$. Then, by writing g as $\langle \alpha, g, g \rangle$ and applying Axiom 4 we have that $f \preceq g$. ∎

From Theorem 2 it is immediate to show that distributions over equivalent alternatives are equivalent themselves. What if some of the alternatives are incomparable? We show that their distributions are either equivalent or incomparable, as due to Axiom 6 no strict preference can be derived.

Lemma 1 *Let $f, g \in \Delta(\mathscr{A})$ be such that, for all $a \in supp(f)$ and $a' \in supp(g)$, $a \not\lesssim a'$. Then, either $f \not\lesssim g$ or $f \sim g$.*

Proof. Let $n_f = |supp(f)|$ and $n_g = |supp(g)|$, we proceed by induction on $n_f + n_g$. If $n_f + n_g = 2$, the thesis is obviously true. Otherwise, assume w.l.o.g. that $n_f > 1$ and let $a \in supp(f)$. We can write f as $\langle f(a), [a], f_{-a} \rangle$, where $f_{-a}(a) = 0$ and $f_{-a}(a') = \frac{f(a')}{1-f(a)}$ for all $a' \in supp(f) \setminus \{a\}$. Since the support of f_{-a} is smaller than the one of f, we can apply the inductive hypothesis to the pair f_{-a}, g, obtaining that either $f_{-a} \not\lesssim g$ or $f_{-a} \sim g$. We can also apply the inductive hypothesis to the pair $[a], g$, obtaining that $[a] \not\lesssim g$ or $[a] \sim g$. Assume by contradiction that $f \bowtie g$ for some $\bowtie \in \{\prec, \succ\}$. By Axiom 6, it holds $[a] \bowtie g$ or $f_{-a} \bowtie g$, which is a contradiction. ∎

Finally, the following generalization of Lemma 1 shows that alternatives that are either incomparable or equivalent lead to distributions that are themselves either incomparable or equivalent.

Theorem 3 *Let $f, g \in \Delta(\mathscr{A})$ be such that, for all $a \in supp(f)$ and $a' \in supp(g)$, either $a \sim a'$ or $a \not\lesssim a'$. Then, $f \not\lesssim g$ or $f \sim g$.*

Proof. The proof is very similar to the one of Lemma 1, where only the base case is affected by the weakened assumption. ∎

5 Conclusions

In this work we challenged the customary decision-theoretic assumption of totality of the preference relations, on the basis of two real-world scenarios. We proposed a weakening of the classical theory and proved that it retains several desirable properties, while allowing for incomparable alternatives.

Partial preferences have been already employed in procurement auctions. In [2] second-price auctions have been generalized by considering a bid domain representing information disclosures and two ad hoc partial preference relations have been defined for modeling the sensitivity of the disclosed data. Then, in [1] the previous framework has been extended for modeling also service cost, QoS and functional differences, etc. Somewhat surprisingly, it has been shown that extending second price auctions to partial preferences does not yield truthful mechanisms, since overbidding may be profitable in some contexts. This means that, in general, partial preferences may significantly change the theoretical properties of a mechanism and the axiomatization presented in this work enables to uniformly employ them in the field of Mechanism Design and estimate their impact.

References

[1] P.A. Bonatti, M. Faella, C. Galdi & L. Sauro (2013): *Auctions for Partial Heterogeneous Preferences*. In: *Proc. of MFCS: Mathematical Foundations of Compuer Science*, *Lecture Notes in Computer Science* 8087, Springer, pp. 183–194, doi:10.1007/978-3-642-40313-2_18.

[2] Piero A. Bonatti, Marco Faella, Clemente Galdi & Luigi Sauro (2011): *Towards a Mechanism for Incentivating Privacy*. In Vijay Atluri & Claudia Díaz, editors: *ESORICS*, *Lecture Notes in Computer Science* 6879, Springer, doi:10.1007/978-3-642-23822-2_26.

[3] R.B. Myerson (1997): *Game Theory: Analysis of Conflict*. Harvard University Press.

[4] L. J. Savage (1954): *The Foundations of Statistics*. John Wiley and Sons, New York.

[5] Herbert Simon (1955): *A Behavioral Model of Rational Choice*. *The Quarterly Journal of Economics* 69(1), pp. 99–118, doi:10.2307/1884852.